

Limits of local-global convergent graph sequences

Hamed Hatami

School of Computer Science, McGill University
Montreal, Quebec, Canada H3A 0G4

László Lovász

Institute of Mathematics, Eötvös Loránd University Budapest, Hungary H-1117

Balázs Szegedy

Department of Mathematics, University of Toronto
Toronto, Ontario, Canada M5S2E4

February 2012

Abstract

The colored neighborhood metric for sparse graphs was introduced by Bollobás and Riordan [8]. The corresponding convergence notion refines a convergence notion introduced by Benjamini and Schramm [6]. We prove that even in this refined sense, the limit of a convergent graph sequence (with uniformly bounded degree) can be represented by a graphing. We study various topics related to this convergence notion such as: Bernoulli graphings, factor of i.i.d. processes and hyperfiniteness.

Contents

| | | |
|-----------|---|-----------|
| 1 | Introduction | 2 |
| 2 | Local-Global convergence of bounded degree graphs | 4 |
| 3 | Involution-invariant measures and graphings | 5 |
| 4 | Local limits of decorated graphs | 8 |
| 5 | A regularization lemma | 9 |
| 6 | Proof of the main theorem | 10 |
| 7 | Bernoulli graphings and Bernoulli graph sequences. | 12 |
| 8 | Non-standard graphings | 16 |
| 9 | Hyperfinite graphs and graphings | 17 |
| 10 | Graphings as operators and expander graphings | 19 |
| 11 | Graphings and local algorithms | 21 |
| 12 | Concluding remarks | 23 |

1 Introduction

The theory of graph convergence is a recently emerging field. It creates a link between combinatorics and analysis similarly as Fürstenberg's correspondence principle connects finite integer sequences with measure preserving systems. Interestingly (or rather unfortunately) there is no unified theory of graph convergence. Instead there are various convergence notions that work well in different situations. For example the theory of dense graph limits [21, 22, 10] works well if the number of edges is quadratic in the number of vertices but it trivializes for sparser graphs. On the other hand the Benjamini–Schramm limit [6] is only defined for graphs which have a linear number of edges in terms of the vertices. In the regime between linear and quadratic the situation is more complicated.

In this paper we focus on the very sparse case where graphs have degrees bounded by some fixed number d (which we consider as fixed throughout). According to Benjamini and Schramm, a graph sequence $(G_n)_{n=1}^\infty$ is convergent if the distribution of the isomorphism types of neighborhoods of radius r (when a vertex is chosen uniformly at random in G_n) converges for every fixed r . This notion of convergence is called *local convergence*, *weak convergence* or *Benjamini–Schramm convergence*.

The following example illustrates why a different, stronger notion of convergence is needed in some cases. For odd n , let G_n be a d -regular expander graph on n nodes. For even n , let G_n be the disjoint union of two d -regular expander graphs on $n/2$ nodes. Assume that the girth of G_n tends to infinity. Then the sequence G_n is locally convergent, but clearly even and odd members

of the sequence are quite different, and it would be desirable to refine our notion of convergence to distinguish them.

Bollobás and Riordan [8] introduced such a finer convergence notion (i.e., fewer sequences are convergent). A graph sequence $(G_n)_{n=1}^\infty$ is convergent in this sense if for every $r, k \in \mathbb{N}$ and $\epsilon > 0$ there is an index l such that if $n, m > l$, then for every coloring of the vertices of G_n with k colors, there is a coloring of the vertices of G_m with k colors such that the distance between the distributions of colored neighborhoods of radius r in G_n and G_m is at most ϵ . This is equivalent to saying that G_n and G_m are close in the colored neighborhood metric introduced in [8]. This finer notion of convergence is sensitive to both local and global properties of the graphs whereas the Benjamini–Schramm convergence is only sensitive to local properties. For this reason we call this notion *local-global convergence*.

Benjamini and Schramm described a limit object for locally convergent sequences in the form of an involution-invariant distribution on rooted countable graphs with bounded degree. One can also describe this limit object as a *graphing* (Elek [14]), which is a bounded degree graph on a Borel probability space such that the edge set is Borel measurable and it satisfies a certain measure preservation property (We will give a precise definition below.). Neighborhood statistics in graphings can be defined by using the probability space structure on the vertex set. Every involution-invariant distribution can be represented by a graphing. We note that graphings are common generalizations of bounded degree graphs and measure preserving systems and so they are also interesting from an ergodic theoretic point of view.

However, the graphing representing the limit object of a locally convergent graph sequence is not unique: different graphings can describe the same involution-invariant distribution. In other words, a graphing contains more information than just the limiting neighborhood distribution. This suggests that graphings can be used to represent limit objects for more refined convergence notions. Indeed, in the present paper we show that the limit of a local-global convergent sequence can also be represented by a graphing in the sense that the graphs in the sequence converge to the graphing in the colored neighborhood metric. This means that for every local-global convergent sequence we produce a graphing which contains both local and global information about the graphs.

We highlight the importance of a special family of graphings called *Bernoulli graphings*. We show that with given local statistics, the Bernoulli graphings contain the least global information. This means that the global properties of a Bernoulli graphing can be modeled with an arbitrary precision on any other graphing with the same local statistics. Graph sequences that converge to Bernoulli graphings will be called *Bernoulli sequences*. For a graph G , being close to a Bernoulli graphing in the colored neighborhood metric, means that the local statistics of a coloring on G can be modeled by a randomized process called *local algorithm* or *factor of i.i.d. process*.

Roughly speaking, a hyperfinite graph sequence is a bounded degree sequence whose members can be cut into small connected components removing a small set of vertices (or equivalently edges). We prove that a locally convergent hyperfinite sequence is a local-global convergent Bernoulli sequence. It is an interesting question how to construct non hyper-finite sequences that are Bernoulli. A good candidate is a growing sequence of random d -regular graphs; however this is a hard open problem. Describing the local-global limits of random d -regular graphs would give a deep understanding of

their structure, and it seems to be one of the most interesting problems in this topic. We describe a few related conjectures.

2 Local-Global convergence of bounded degree graphs

A *rooted graph* is a pair (G, o) where o is a vertex of a graph G . The *radius* of a rooted graph is the distance of the farthest vertex in G to o . We denote by U^r the set of all rooted graphs with radius at most r (and all degrees bounded by d). For an integer $r \geq 0$, and a vertex v in a graph G , let $N_{G,r}(v)$ denote the subgraph of G rooted at v and induced by the vertices that are at a distance at most r from v . Two rooted graphs (G, o) and (G', o') are said to be *isomorphic* if there is an isomorphism from G to G' that maps o to o' .

Given a finite graph G and a radius $r \geq 0$, we can choose a node $v \in V(G)$ uniformly and randomly, and consider the distribution of $N_{G,r}(v)$. Let $P_{G,r}$ denote this probability measure on U^r . We say that a sequence (G_n) of finite graphs is *locally convergent* (or Benjamini–Schramm convergent) if $P_{G_n,r}$ converges to a limit distribution as $n \rightarrow \infty$, for every fixed $r \geq 0$.

Denote the set of probability measures on a Borel space X by $M(X)$. Note that since U^r is finite, all the usual distances on $M(U^r)$ are topologically equivalent. We shall usually work with the *total variation distance* d_{var} , defined (in general, for a space X) by

$$d_{\text{var}}(\mu, \nu) = \sup_{A \subseteq X} |\mu(A) - \nu(A)|$$

where A runs through the Borel measurable sets.

To define our refinement of local convergence, we consider vertex colorings. For a finite graph G , let $K(k, G)$ denote the set of all vertex colorings with k colors. Fix integers k and r , and let $U^{r,k}$ be the set of all triples (H, o, c) where (H, o) is a rooted graph of radius at most r and c is an arbitrary k -coloring of $V(H)$. Consider a finite graph G together with a $c \in K(k, G)$. Pick a random vertex v from G . Then the restriction of the k -coloring to $N_{G,r}(v)$ is an element in $U^{r,k}$, and thus for the graph G , every $c \in K(k, G)$ introduces a probability distribution on $U^{r,k}$ which we denote by $P_{G,r}[c]$. Sometimes we refer to the probability distributions $P_{G,r}[c]$ (for $r \geq 0$) as *local statistics* of the coloring c . Let

$$Q_{G,r,k} := \{P_{G,r}[c] : c \in K(k, G)\} \subseteq M(U^{r,k}).$$

These sets are similar to “quotient sets” introduced in [11] for dense graphs, except that there only edges with the given coloring were counted, while here we consider the colors on larger neighborhoods. Notice that the sets $Q_{G,r,k}$ are finite, and they are subsets of the finite dimensional space $\mathbb{R}^{U^{r,k}}$ that is independent of the graph G .

Definition 2.1 A sequence of finite graphs $(G_n)_{n=1}^\infty$ with all degrees at most d is called *locally-globally convergent* if for every $r, k \geq 1$, the sequence $(Q_{G_n,r,k})_{n=1}^\infty$ converges in the Hausdorff distance inside the compact metric space $(M(U^{r,k}), d_{\text{var}})$.

Since compact subsets of a compact metric space form a compact space with respect to the Hausdorff metric, it follows that every infinite sequence of finite graphs contains a locally-globally convergent subsequence.

Fixing $k = 1$ in Definition 2.1, we recover a metric definition of Benjamini–Schramm convergence. It is easy to construct examples of graph sequences which are convergent in Benjamini–Schramm sense, but not locally-globally. However we do not know whether $k = 2$ would give a convergence notion equivalent to local-global convergence.

It is natural to ask if we obtain a different convergence notion if we replace vertex colorings by edge colorings or other locally defined extra structures. It turns out that all local structures can be encoded by vertex colorings, and thus they do not lead to different convergence notions. As an example, we show how to encode edge colorings by vertex colorings.

Let G be a graph with all degrees at most d and let $c : E(G) \rightarrow [k]$ be an edge coloring of G . It is easy to see that there exists an edge coloring $c_1 : E(G) \rightarrow [30d^3k]$ such that $c_1(e) \equiv c(e)$ modulo k for every $e \in E(G)$, and if $c_1(e_1) = c_1(e_2)$ holds, then the edges e_1 and e_2 are of distance at least 3 in the edge graph of G . It is clear that c_1 encodes the coloring c in the sense that local statistics of c_1 modulo k give the local statistics of c . Let S denote the set of subsets of $[30d^3k]$ of size at most d . We define the vertex coloring $c_2 : V(G) \rightarrow S$ by setting $c_2(v)$ to be the set of c_1 -colors of the edges incident to v . Now it is easy to see that c_2 encodes the coloring c_1 in the following way. If $e = (v, w)$ is an edge in G , then $c_1(e)$ is the intersection of the sets $c_2(v)$ and $c_2(w)$.

3 Involution-invariant measures and graphings

Benjamini and Schramm [6] associated a limit object with every locally convergent graph sequence as follows. Let \mathfrak{G} denote the set of (isomorphism classes of) rooted, connected (possibly infinite) graphs with all degrees at most d . For a rooted graph (B, o) with radius r , we denote by $\mathfrak{G}(B, o)$ the set of all rooted graphs (G, o) such that $N_{G,r}(o) \cong (B, o)$. For a rooted graph (G, o) , we define a neighborhood basis at (G, o) as $\mathfrak{G}(N_{G,r}(o))$. These neighborhoods define a topology on \mathfrak{G} . It is easy to see that this is a compact separable space.

The Benjamini–Schramm limit of the locally convergent graph sequence $(G_n)_{n=1}^\infty$ is a probability measure ν on the Borel sets of \mathfrak{G} , such that

$$\lim_{n \rightarrow \infty} P_{G_n, r}(B, o) = \nu(\mathfrak{G}(B, o))$$

for every $r \geq 1$ and every rooted graph (B, o) of radius r .

Not every probability measure on \mathfrak{G} arises as the limit of a convergent graph sequence. One property that all limits have is called *involution invariance* or *unimodularity*. To define this, let $\tilde{\mathfrak{G}}$ denote the space of graphs in \mathfrak{G} with a distinguished edge incident to the root. Let $\alpha : \tilde{\mathfrak{G}} \rightarrow \tilde{\mathfrak{G}}$ denote the continuous transformation that moves the root to the other endpoint of the distinguished edge. For every probability measure μ on \mathfrak{G} , define μ^* to be the unique probability measure on \mathfrak{G} such that $d\mu^*/d\mu(G)$ is proportional to the degree of the root in G . Define the probability measure $\tilde{\mu}$ on $\tilde{\mathfrak{G}}$ by first picking a μ^* -random graph, and then distinguishing a random edge incident to the root. The measure μ is called *involution-invariant* if $\tilde{\mu}$ is invariant under α . Involution-invariant measures on \mathfrak{G} form a closed set in the weak topology.

Let G be a finite graph, and let the probability measure ν on the Borel sets of \mathfrak{G} be defined as $\nu(\mathfrak{G}(B, o)) = P_{G,r}(B, o)$ for every $r \geq 1$ and every rooted graph (B, o) of radius r . It is easy to see

that ν is involution-invariant. It follows that every measure on \mathfrak{G} that is the limit of finite graphs is involution-invariant. Aldous and Lyons [1] conjectured that all involution-invariant measures arise as graph limits. The Aldous-Lyons conjecture is considered to be one of the most important open problems in this area.

In the dense setting, the set of the symmetric measurable maps $w : [0, 1]^2 \rightarrow [0, 1]$ were used to generalize the concept of graphs and describe graph limits [21]. For local-global convergence of Definition 2.1, graphings serve this purpose.

Definition 3.1 Let X be a Polish topological space and let ν be a probability measure on the Borel sets in X . A *graphing* is a graph \mathcal{G} on $V(\mathcal{G}) = X$ with Borel measurable edge set $E(\mathcal{G}) \subset X \times X$ in which all degrees are at most d and

$$\int_A e(x, B) d\nu(x) = \int_B e(x, A) d\nu(x) \quad (1)$$

for all measurable sets $A, B \subseteq X$, where $e(x, S)$ is the number of edges from $x \in X$ to $S \subseteq X$.

Note that every finite graph G is a graphing where $X = V(G)$ and ν_G is the uniform distribution on $V(G)$.

If (1) holds, then $\eta^*(A \times B) = \int_A e(x, B) d\nu(x)$ defines a measure on the Borel sets of $X \times X$. This measure is concentrated on $E(\mathcal{G})$, symmetric in the two coordinates, and its marginal ν^* satisfies $(d\nu^*/d\nu)(x) = \deg(x)$. Normalizing by $d_0 = \int_X \deg(x) dx$, we get a probability distribution η on the set of edges. We can generate a random edge from η by selecting a random point v from ν^* and selecting uniformly a random edge incident with v . Conversely if we generate a random oriented edge this way, and the distribution that is obtained is invariant under flipping the orientation, then (1) follows by Fubini's theorem.

Let \mathcal{G} be a graphing (of degree at most d) on the probability space (X, ν) . Then it induces a measure $\mu_{\mathcal{G}}$ on \mathfrak{G} : pick a random element $x \in X$ and take its connected component \mathcal{G}_x rooted at x . It is easy to see that $\mu_{\mathcal{G}}$ is an involution-invariant measure (In fact, (1) just expresses this property.).

Let \mathcal{G} be a graphing as in Definition 3.1. A *vertex coloring* of \mathcal{G} with k colors is a measurable function $c : X \rightarrow [k]$. The set of all such colorings will be denoted by $K(k, \mathcal{G})$. We define $P_{\mathcal{G}, r}[c]$ and $Q_{\mathcal{G}, r, k}$ in a similar way as in a finite graph (Notice that it makes sense to talk about a random vertex in \mathcal{G}). The set $Q_{\mathcal{G}, r, k}$ is a subset of the finite dimensional space $\mathbb{R}^{U^{r, k}}$, but in general it might be infinite and not necessarily closed; we will often use its closure $\overline{Q}_{\mathcal{G}, r, k}$ (see Question 8.1).

Now we are ready to state our main theorem.

Theorem 3.2 *Let $(G_i)_{i=1}^\infty$ be a local-global convergent sequence of finite graphs with all degrees at most d . Then there exists a graphing \mathcal{G} such that $Q_{G_n, r, k} \rightarrow Q_{\mathcal{G}, r, k}$ ($n \rightarrow \infty$) in Hausdorff distance for every r and k .*

To what degree is the limit object determined? This question leads to different notions of “isomorphism” between graphings.

Definition 3.3 Let $(\mathcal{G}_1, X_1, \nu_1)$ and $(\mathcal{G}_2, X_2, \nu_2)$ be graphings.

- They are called *locally equivalent* if for every $r \in \mathbb{N}$, the distribution of $N_{\mathcal{G}_1, r}(x_1)$ is the same as the distribution of $N_{\mathcal{G}_2, r}(x_2)$ for random $x_1 \in X_1$ and $x_2 \in X_2$.
- They are called *locally-globally equivalent* if $\overline{Q}_{\mathcal{G}_1, r, k} = \overline{Q}_{\mathcal{G}_2, r, k}$ for every $r, k \in \mathbb{N}$.

Local equivalence of two graphings means that they induce the same involution-invariant measure on \mathfrak{G} . Local-global equivalence implies local equivalence by setting $k = 1$.

Definition 3.4 (Local-global partial order) Assume that \mathcal{G}_1 and \mathcal{G}_2 are two graphings of maximal degree at most d . We say that $\mathcal{G}_1 \prec \mathcal{G}_2$ if $\overline{Q}_{\mathcal{G}_1, r, k} \subseteq \overline{Q}_{\mathcal{G}_2, r, k}$ for every $r, k \geq 1$. In particular, \mathcal{G}_1 and \mathcal{G}_2 are locally-globally equivalent if and only if both $\mathcal{G}_1 \prec \mathcal{G}_2$ and $\mathcal{G}_2 \prec \mathcal{G}_1$ hold.

Recall that a measurable map $\phi : (X, \mu) \rightarrow (Y, \nu)$ is called measure-preserving if $\mu(\phi^{-1}(A)) = \nu(A)$ for every measurable set $A \subseteq Y$. An easy way to prove a relation $\mathcal{G}_1 \prec \mathcal{G}_2$ between two graphings is the following. We call a measure preserving map $\phi : V(\mathcal{G}_1) \rightarrow V(\mathcal{G}_2)$ a *local isomorphism* if restricted to any connected component of \mathcal{G}_1 , we get an isomorphism with a connected components of \mathcal{G}_2 . Clearly local isomorphisms can be combined. However, a local isomorphism may not be invertible! It is easy to see that the existence of a local isomorphism $\mathcal{G}_1 \rightarrow \mathcal{G}_2$ implies that \mathcal{G}_1 and \mathcal{G}_2 are locally equivalent, and $\mathcal{G}_2 \prec \mathcal{G}_1$.

Example 3.5 Let G be a finite connected graph, and $G \cup G$ denote the disjoint union of G with itself. The function $\phi : V(G \cup G) \rightarrow V(G)$ that maps both copies of G in $G \cup G$ isomorphically to G is a (non-invertible) local isomorphism. Consequently $G \cup G$ and G are locally equivalent, and $G \prec G \cup G$. Note however that G and $G \cup G$ are not locally-globally equivalent.

We shall study the local-global equivalence and the local-global partial order in Section 7. In particular, we will show that among all graphings in a local equivalence class, there is always a smallest and a largest in this partial order.

We conclude this section with a few remarks.

Remark 3.6 (Directed graphings) Let X and ν be as in Definition 3.1 and let $E(\mathcal{G}) \subset X \times X$ be the edge set of a directed Borel graph \mathcal{G} of bounded degree. For two set $A, B \subset X$, let $e(A, B) = |E(\mathcal{G}) \cap A \times B|$ denote the number of directed edges from A to B (This quantity may be infinite.). Then \mathcal{G} is called a (directed) graphing if

$$\int_A e(x, B) \, d\nu = \int_B e(A, x) \, d\nu$$

holds for any two measurable sets A, B .

Remark 3.7 (Some examples) The following simple example for a directed graphing comes from ergodic theory. Let $T : X \rightarrow X$ be a measure preserving transformation which has a measure preserving inverse. Then the graph $\{(x, T(x)) | x \in X\}$ is a directed graphing. More specifically, let θ be an irrational number and $T(x) = x + \theta$ on the circle group \mathbb{R}/\mathbb{Z} . Then $\{(x, x + \theta) | x \in \mathbb{R}/\mathbb{Z}\}$ is an “ergodic” directed graphing (Ergodic graphings are defined in Section 10.). A similar but undirected graphing on \mathbb{R}/\mathbb{Z} is given by the edge set $\{(x \pm \theta) | x \in \mathbb{R}/\mathbb{Z}\}$.

Remark 3.8 (Decomposition into maps) The following construction can be used to verify the graphing axiom (1) in some cases. Let (X, ν) be a measure space. A *measure preserving equivalence* between two measurable sets $A, B \subset X$ is a map $\psi : A \rightarrow B$ which is measure preserving and has a measure preserving inverse. A partial measure preserving equivalence between A and B is a measure preserving equivalence between $A' \subset A$ and $B' \subset B$. Let $X = X_1 \cup X_2 \cup \dots \cup X_n$ be an almost disjoint (intersections have 0 measure) decomposition of X and for each pair $1 \leq i < j \leq n$, let $\psi_{i,j}$ be a partial measure preserving equivalence between X_i and X_j . Then the symmetrized version of $E = \cup_{1 \leq i < j \leq n} \{(x, \psi_{i,j}(x)) | x \in X_i\}$ is the edge set of a graphing. It is not hard to prove that each graphing has such a decomposition. In fact any measurable coloring of the vertex set in which vertices of the same color are of distance at least 3 yields such a decomposition. The existence of such a coloring follows from results of Kechris, Solecki and Todorcevic [19]. This construction gives an upper bound of $O(d^2)$ for the number of maps. With more care (considering the line graph of \mathcal{G}), one can reduce this to $2d - 1$. It is not known whether $d + 1$ maps would suffice. We come back to this issue later in Section 12.

4 Local limits of decorated graphs

In this section we extend the formalism behind the Benjamini–Schramm limits for the case when vertices are decorated by elements from a compact space. Let C be a second countable compact Hausdorff space. Let $\mathfrak{G}(C)$ denote the space of (isomorphism classes of) rooted, connected (countable) graphs with all degrees at most d such that the vertices are decorated by elements from C . So the points of $\mathfrak{G}(C)$ are triples (G, o, c) , where G is a connected countable graph, $o \in V(G)$, and $c : V(G) \rightarrow C$. $c : V(G) \rightarrow C$ is a feature of the graph G . If C is the trivial (one point) compact space, then $\mathfrak{G}(C)$ can be identified with the space \mathfrak{G} defined earlier. Two important special cases for us will be when $C = [0, 1]$ (assigning $[0, 1]$ weights to vertices), and $C = [k]$ (coloring vertices by k colors). With a slight abuse of notation, these will be denoted by $\mathfrak{G}[0, 1]$ and $\mathfrak{G}[k]$.

We put a compact topology on $\mathfrak{G}(C)$ by specifying a basis of it. Let r be an arbitrary natural number and (H, o) be a finite rooted graph of radius r . Assume furthermore that every vertex v of (H, o) is decorated by an open set U_v in C . Let S be the collection of all $(G, o, c) \in \mathfrak{G}(C)$ where the neighborhood $N_{G,r}(o)$ is isomorphic to (H, o) , and furthermore there is an isomorphism $\alpha : N_{G,r}(o) \rightarrow (H, o)$ such that $c(v) \in U_{\alpha(v)}$ for every $v \in N_{G,r}(o)$. It is easy to see that $\mathfrak{G}(C)$ with this topology is a compact, second countable, Hausdorff space. As a consequence, probability measures on $\mathfrak{G}(C)$ form a compact space in the weak topology.

Let G be a finite graph with all degrees at most d in which the vertices are C -labeled. We can construct a probability measure μ_G on $\mathfrak{G}(C)$ by putting a root o on a randomly chosen vertex $v \in V(G)$ and keeping only the connected component of the root. A sequence $(G_n)_{n=1}^\infty$ of C -labeled graphs is called *locally convergent* if the corresponding measures $\{\mu_{G_n}\}_{n=1}^\infty$ converge in the weak topology to some measure μ . The measure μ is the limit object of the sequence.

We define involution-invariance completely analogously to the undecorated case, simply replacing \mathfrak{G} by $\mathfrak{G}(C)$ everywhere. Involution-invariant measures on $\mathfrak{G}(C)$ form a closed set in the weak topology. It follows that if μ is a measure on $\mathfrak{G}(C)$ that is the limit of finite C -decorated graphs,

then it is involution-invariant.

A *C-decorated graphing* is a graphing \mathcal{G} together with a Borel function $c : V(\mathcal{G}) \rightarrow C$. Similarly as in the undecorated case, every *C-decorated graphing* defines an involution-invariant distribution. The measure $\mu_{\mathcal{G},c}$ on $\mathfrak{G}(C)$ is created by picking a random element $x \in V(\mathcal{G})$, and taking its connected component \mathcal{G}_x rooted at x together with the vertex labels given by the restriction of c to $V(\mathcal{G}_x)$. It is easy to see that $\mu_{\mathcal{G},c}$ is an involution-invariant measure.

Remark 4.1 We can define a Borel graph on $\mathfrak{G}(C)$. The edge set $\mathcal{E}(C)$ of this graph consists of pairs $((G, o_1, c), (G, o_2, c)) \in \mathfrak{G}(C) \times \mathfrak{G}(C)$ such that (o_1, o_2) is an edge in G . Note that loop edges can arise in this graph. For example if there is an automorphism of (G, c) which takes o_1 to its neighbor o_2 , then (G, o_1, c) is identified with (G, o_2, c) in $\mathfrak{G}(C)$. In general it is not true that every involution-invariant measure ν on $\mathfrak{G}(C)$ turns this graph into a graphing. This is due to the problem with automorphisms which also lead to loops. However it is not hard to show that if for an involution-invariant measure ν , with probability one, a ν -random connected component has no automorphisms, then we get a graphing $(\mathfrak{G}(C), \nu, \mathcal{E}(C))$. One important role of appropriate decorations is to break symmetries, and make this graph a graphing.

Definition 4.2 Let $f : \mathfrak{G}(C) \rightarrow \mathbb{R}$ be any function. We call f *local* if there is a positive integer r such that $f(G, o, c)$ depends (measurably) on the decorated r -neighborhood $N_{G,r}(o)$.

Clearly every local function is measurable. Local functions correspond to those functions on finite graphs that can be computed by a constant-time local algorithm. We shall discuss this in more details in Section 11.

5 A regularization lemma

The following lemma is the main ingredient in proving Theorem 3.2. It serves as a “regularity lemma” in our framework for bounded degree graphs.

Lemma 5.1 (Regularization) *For positive integers r, k and real number $\epsilon > 0$, there exists an integer $t_{r,k,\epsilon}$ such that the following holds. For every graph G with all degrees at most d , there exists a $t_{r,k,\epsilon}$ -vertex coloring q of G which satisfies the following conditions.*

- *If $q(v) = q(w)$, then either $v = w$ or the distance of v and w in G is at least $r + 1$;*
- *For every $g \in K(k, G)$, there exists $\alpha : [t_{r,k,\epsilon}] \rightarrow [k]$ such that*

$$d_{\text{var}}(P_{G,r}[g], P_{G,r}[\alpha \circ q]) \leq \epsilon.$$

Proof. The space $M(U^{r,k})$ is a bounded dimensional compact set with the topology generated by d_{var} . Let N be an $\epsilon/2$ -net in $M(U^{r,k})$ in d_{var} . Let N_G be the subset of points in N that are at most $\epsilon/2$ far from a point of the form $P_{G,r}[g]$ for some $g \in K(k, G)$. For each $a \in N_G$, we choose a representative $x_a = P_{G,r}[g_a]$ such that $d_{\text{var}}(a, x_a) \leq \epsilon/2$. It is clear that for every $g \in K(k, G)$, there is a point x_a such that $d_{\text{var}}(P_{G,r}[g], x_a) \leq \epsilon$. Let f be the common refinement of all the partitions

$\{g_a\}_{a \in N_G}$. Clearly f has a bounded number of partition sets in terms of r, k, ϵ and d and it satisfies the second condition.

Now we further refine f to satisfy the first condition. Let f' be a proper coloring of the graph G with $(d+1)^r$ colors in which every two vertices in distance at most r receive different colors. The common refinement q of f and f' satisfies both conditions. ■

6 Proof of the main theorem

Now we introduce the space X which will serve as a universal Borel space for the limit graphings of sequences of finite graphs with all degrees at most d . Consider the compact space $C = \prod_{k,r,n} [t_{r,k,1/n}]$ with the product topology where $t_{r,k,1/n}$ are defined according to Lemma 5.1. We denote by X the compact space $\mathfrak{G}(C)$ and by $E \subset X \times X$ the set of edges $((G, o_1, c), (G, o_2, c))$ such that (o_1, o_2) is an edge in G (See Remark 4.1). Let $q : X \rightarrow C$ be the function defined as $q : (G, o, c) \mapsto c(o)$. Furthermore for $r, k, n \in \mathbb{N}$, define the coloring $q_{r,k,n} : X \rightarrow [t_{r,k,1/n}]$ as the composition of q with the projection to the coordinate (r, k, n) in C .

Let $(G_i)_{i=1}^\infty$ be a local-global convergent sequence of graphs with all degrees at most d . For each G_i and triple $(r, k, n) \in \mathbb{N}^3$, we choose a coloring $q_{r,k,n}^i : V(G_i) \rightarrow [t_{r,k,1/n}]$ guaranteed by Lemma 5.1. Let $q_i : V(G_i) \rightarrow C$ be defined as $\prod_{r,k,n} \{q_{r,k,n}^i(v)\} \in C$. As described in Section 4, each graph G_i together with the coloring q_i defines a probability measure μ_i on X by putting the root on a random vertex of G_i and keeping only the connected component of the root.

By choosing a subsequence from $(G_i)_{i=1}^\infty$ we can assume that the sequence $\{\mu_i\}_{i=1}^\infty$ weakly converges to a probability distribution μ on X . Our goal is to show that the Borel graph (X, E) with the measure μ is a graphing which represents the local-global limit of $(G_i)_{i=1}^\infty$.

Let us first observe that for a μ -random element (G, o, c) in (X, μ) , with probability one, the vertex labels $\{c(v) : v \in V(G)\}$ are all different. This follows from the fact that the colorings $q_{r,k,n}^i$ separate points in G_i that are closer than $r+1$, and that this property is preserved in the limit. This means that if $v, w \in V(G)$ are of distance r , then with probability one their colors projected to the coordinate (r, k, n) (where k, n are arbitrary) are different.

Lemma 6.1 *The measurable graph (X, E, μ) is a graphing.*

Proof. Let us introduce the measures $\{\eta_i^*\}_{i=1}^\infty$, similarly as in Section 3, by

$$\eta_i^*(A \times B) = \int_A e(x, B) d\mu_i(x),$$

where $A, B \subseteq X$ are measurable, and $e(x, B)$ is the number of edges $(x, y) \in E$ with $y \in B$. We define η^* analogously as $\eta^*(A \times B) = \int_A e(x, B) d\mu(x)$.

Assume that $A, B \subset X$ are open-closed sets. The weak convergence of $\{\mu_i\}_{i=1}^\infty$ implies that $\lim_{i \rightarrow \infty} \eta_i^*(A \times B) = \eta^*(A \times B)$ and $\lim_{i \rightarrow \infty} \eta_i^*(B \times A) = \eta^*(B \times A)$. Note that $\eta_i^*(A \times B) = \eta_i^*(B \times A)$, since both are equal (up to normalization by $|V(G_i)|$) to the number of edges between the sets $\{v | (G_i, v, q_i) \in A\}$ and $\{v | (G_i, v, q_i) \in B\}$. Here we used the fact that the vertex labels $q_i(\cdot)$ are all different and thus automorphisms of G_i cannot cause any problems. We obtain that

$\eta^*(B \times A) = \eta^*(A \times B)$, and since such product sets generate the whole σ -algebra on $X \times X$, the proof is complete. \blacksquare

Lemma 6.2 *The probability distributions $P_{G_{i,r}}[q_{r,k,n}^i]$ converge to $P_{\mathcal{G},r}[q_{r,k,n}]$ for every fixed triple $r, k, n \in \mathbb{N}$.*

Proof. Pick a μ -random point $x = (G, o, c) \in X$. Let the rooted graph \mathcal{G}_x be the connected component of x in the graphing \mathcal{G} rooted at x . There is a natural vertex coloring on \mathcal{G}_x which is the restriction of the function q to the vertices of \mathcal{G}_x . So \mathcal{G}_x can be regarded as an element in X . We claim that with probability one $x = (G, o, c)$ is isomorphic (in a root and label preserving way) to $(\mathcal{G}_x, q|_{\mathcal{G}_x})$. Indeed with probability one all the vertex labels of G are different, and in this case the map given by $v \mapsto (G, v, c)$ defines a decoration-preserving isomorphism between (G, o, c) and \mathcal{G}_x (The fact that the vertex labels in G are all different guarantees that the map is one to one.).

We conclude that the probability distribution $P_{\mathcal{G},r}[q_{r,k,n}]$ is the same as the distribution of $(N_{G,r}(o), c_{r,k,n})$ where (G, o, c) is a μ -random element in X , and $c_{r,k,n}$ is the projection of c to the coordinate (r, k, n) . The lemma now follows from the weak convergence of $\{\mu_i\}_{i=1}^\infty$ to μ . \blacksquare

Lemma 6.3 *For every $r, k, n \in \mathbb{N}$, there is an index i_0 such that for every $i \geq i_0$ and $c \in K(k, G_i)$, there is a k -coloring c' of X such that $d_{\text{var}}(P_{G_i,r}[c], P_{\mathcal{G},r}[c']) \leq 1/n$.*

Proof. By Lemma 6.2 there is an index i_0 such that

$$d_{\text{var}}(P_{G_{i,r}}[q_{r,k,2n}^i], P_{\mathcal{G},r}[q_{r,k,2n}]) \leq \frac{1}{2n} \quad (2)$$

for every index $i \geq i_0$. Let $i \geq i_0$ be arbitrary, and let $c \in K(k, G_i)$ be a k -coloring of G_i . Then by Lemma 5.1 there is a map $\alpha : [t_{r,k,1/(2n)}] \rightarrow [k]$ such that

$$d_{\text{var}}(P_{G_{i,r}}[\alpha \circ q_{r,k,2n}^i], P_{G_{i,r}}[c]) \leq \frac{1}{2n}.$$

On the other hand the definition of the total variation distance and (2) guarantee that

$$d_{\text{var}}(P_{G_{i,r}}[\alpha \circ q_{r,k,2n}^i], P_{\mathcal{G},r}[\alpha \circ q_{r,k,2n}]) \leq \frac{1}{2n}.$$

Hence $c' = \alpha \circ q_{r,k,2n}$ satisfies the required condition. \blacksquare

Lemma 6.4 *For every coloring $c \in K(k, \mathcal{G})$ and natural numbers n, r , there is an index i_0 such that for every $i \geq i_0$, there is a coloring $c' \in K(k, G_i)$ with $d_{\text{var}}(P_{G_{i,r}}[c'], P_{\mathcal{G},r}[c]) \leq 1/n$.*

Proof. Let $c : X \rightarrow [k]$ be a Borel coloring. Then for every $\epsilon > 0$, there is a continuous coloring $c_\epsilon : X \rightarrow [k]$ such that $|\mu(c^{-1}(a) \triangle c_\epsilon^{-1}(a))| \leq \epsilon$ for all $1 \leq a \leq k$. Taking ϵ to be sufficiently small, we have

$$d_{\text{var}}(P_{\mathcal{G},r}[c_\epsilon], P_{\mathcal{G},r}[c]) \leq \frac{1}{2n}. \quad (3)$$

Let the graphing \mathcal{G}_i be the same as the graphing \mathcal{G} with the only difference that the measure μ is replaced by μ_i . Since $\{\mu_i\}_{i=1}^\infty$ converges weakly to μ and c_ϵ is continuous, there is an index i_0 such that if $i \geq i_0$, then

$$d_{\text{var}}(P_{\mathcal{G}_i,r}[c_\epsilon], P_{\mathcal{G},r}[c_\epsilon]) \leq \frac{1}{2n}. \quad (4)$$

The coloring c_ϵ induces a coloring f_ϵ^i on G_i which assigns to every vertex $v \in V(G_i)$ the c_ϵ color of the rooted graph $(G_i, v, q_i) \in X$. Then we have $P_{G_i, r}[f_\epsilon^i] \equiv P_{\mathcal{G}_i, r}[c_\epsilon]$ which together with (3) and (4) completes the proof. \blacksquare

7 Bernoulli graphings and Bernoulli graph sequences.

Probably the most fundamental graphing construction is the Bernoulli graphing corresponding to an involution-invariant measure. These graphings are closely related to factor of i.i.d. processes and local algorithms. In this chapter we explain their role in local-global convergence.

Definition 7.1 (Bernoulli graphings) Let μ be an involution-invariant measure on \mathfrak{G} . Let ν be the probability measure on $\mathfrak{G}[0, 1]$ produced by putting independent random elements from $[0, 1]$ on the nodes of a μ -random graph. The triple $(\mathfrak{G}[0, 1], \nu, \mathcal{E}[0, 1])$ as defined in Remark 4.1 will be called the *Bernoulli graphing* corresponding to μ , and denoted by \mathcal{B}_μ .

It is not hard to see that \mathcal{B}_μ is a graphing and it represents the involution-invariant distribution μ (Elek [14]).

Remark 7.2 Perhaps it would be more natural to decorate the nodes of the μ -random graph by independent bits, or more generally, by colors from $[k]$ for some fixed $k \geq 2$. This would yield an involution-invariant distribution on $\mathfrak{G}[k]$, but the graph $(\mathfrak{G}[k], \mathcal{E}[k])$ together with this distribution would not necessarily form a graphing.

We define the Bernoulli graphing $\mathcal{B}_\mathcal{G}$ corresponding to an arbitrary graphing \mathcal{G} as the Bernoulli graphing defined by the involution-invariant distribution induced by \mathcal{G} on \mathfrak{G} . Clearly \mathcal{G} and $\mathcal{B}_\mathcal{G}$ are locally equivalent.

Example 7.3 A simple example for a Bernoulli graphing is provided by the involution-invariant measure which is concentrated on a single d -regular rooted tree. Let T denote the rooted d -regular tree, and let (X, ν) be the probability space in which we put independent random labels from $[0, 1]$ on the vertices of T . Two points of X are connected in \mathcal{G} if they can be obtained from each other by moving the root to a neighboring vertex. It seems to be an interesting problem to decide whether the sets $Q_{\mathcal{G}, r, k}$ are all closed (see also Question 8.1).

The following is a related construction. For every graphing \mathcal{G} on the probability space (X, ν) , we define its *Bernoulli lift* \mathcal{G}^+ as follows. The underlying set X^+ of \mathcal{G}^+ will be pairs (x, ξ) , where $x \in X$ and $\xi : V(\mathcal{G}_x) \rightarrow [0, 1]$ assigns labels from $[0, 1]$ to the vertices of the connected component \mathcal{G}_x rooted at x . We connect (x, ξ) to (y, ν) if y is a neighbor of x and $\xi = \nu$ (Note that if y is a neighbor of x , then $\mathcal{G}_x = \mathcal{G}_y$). The measure on X^+ is defined as follows. To generate a random element of X^+ , one picks a ν -random point $x \in X$, and then assigns independent random weights $\xi(u)$ to the nodes u of \mathcal{G}_x .

We define two maps $\phi : V(\mathcal{G}^+) \rightarrow V(\mathcal{G})$ and $\psi : V(\mathcal{G}^+) \rightarrow V(\mathcal{B}_\mathcal{G})$ by $\phi(x, \xi) = x$ and $\psi(x, \xi) = (\mathcal{G}_x, \xi)$. It is easy to check that the maps ϕ and ψ are local isomorphisms. This implies that graphing \mathcal{G} is locally equivalent to its Bernoulli lift \mathcal{G}^+ as well as its Bernoulli graphing $\mathcal{B}_\mathcal{G}$.

Our main goal in this section is to describe the relationship between \mathcal{G} , $\mathcal{B}_{\mathcal{G}}$ and \mathcal{G}^+ from the point of view of local-global equivalence.

Definition 7.4 A graphing is called *atom-free* if its underlying probability space contains no mass points.

Remark 7.5 Note that no finite graph corresponds to an atom-free graphing. Using the graphing property (1) it is easy to see that if a graphing contains an atom, then this belongs to a finite component. If \mathcal{G} is the local limit of a sequence of connected graphs $(G_n)_{n=1}^{\infty}$ with $|V(G_n)| \rightarrow \infty$, then all its components are infinite, and hence it is atom-free. On the other hand, if the union of finite components of a graphing has positive weight, then merging isomorphic finite components we get atoms. Furthermore if \mathcal{G} is the local-global limit of graphs $(G_n)_{n=1}^{\infty}$ (not necessarily connected) with $|V(G_n)| \rightarrow \infty$, then \mathcal{G} is atom-free. This follows from the observation that a graphing is atom-free if and only if its points have a Borel k -coloring with equal color classes for every k .

The following is our main result in this section.

Theorem 7.6 *Every atom-free graphing is local-global equivalent to its Bernoulli lift.*

The map $\psi : V(\mathcal{G}^+) \rightarrow V(\mathcal{B}_{\mathcal{G}})$ defined above is a local isomorphism from \mathcal{G}^+ to $\mathcal{B}_{\mathcal{G}}$. Thus we have the relation $\mathcal{B}_{\mathcal{G}} \prec \mathcal{G}^+$ which implies as a corollary to Theorem 7.6 that Bernoulli graphings are minimal elements in the set of atom-free graphings in their local equivalence class. A group theoretical analogue of this fact was obtained in [3].

Corollary 7.7 (Minimality of Bernoulli graphings) *For every atom-free graphing \mathcal{G} , we have $\mathcal{B}_{\mathcal{G}} \prec \mathcal{G}$.*

In an algorithmic setting, a Borel coloring of \mathcal{G}^+ can be considered as a coloring that depends not only on the graph, but also on a random real number at each point. To be able to imitate this in \mathcal{G} , we have to construct “random-like” colorings on \mathcal{G} . For technical reasons, we have to deal with graphings that already have a Borel coloring.

Definition 7.8 (Quasirandom colorings) Let \mathcal{G} be a graphing on the space (X, ν) , and let $h : X \rightarrow [l]$ be a Borel coloring. Let $\mu_{r,h,k}$ be the probability distribution on $U^{r,kl}$ obtained from ν by considering the r -neighborhood of a random element $x \in X$ and decorating its vertices by random independent elements from $[k]$ (in addition to the given l -coloring h). We say that a measurable coloring $c : X \rightarrow [k]$ is (r, ϵ) -*quasirandom* if $d_{\text{var}}(P_{\mathcal{G},r}[c \times h], \mu_{r,h,k}) \leq \epsilon$ where $c \times h$ denotes the kl -coloring with pairs of colors $(c(x), h(x))$.

Lemma 7.9 (Existence of quasirandom colorings) *Let \mathcal{G} be a atom-free graphing on the space (X, ν) . Then for every $k, r, l \in \mathbb{N}$, $\epsilon > 0$ and Borel l -coloring h , there is an (r, ϵ) -quasirandom coloring $c : X \rightarrow [k]$ of (\mathcal{G}, h) .*

Proof. Let $C = \{0, 1\}^{\mathbb{N}}$ be the Cantor set with the uniform measure. Since (X, ν) has no mass points, there is a measurable equivalence between C and X , without loss of generality, we can

identify the two spaces, assume that $X = C$. Let $\pi_i : C \rightarrow \{0, 1\}^i$ be the projection to the first i coordinates. The map π_i is measure preserving if we consider the uniform measure on $\{0, 1\}^i$. Fix $k, r \in \mathbb{N}$, and let $g_i : \{0, 1\}^i \rightarrow [k]$ be a uniform random coloring of $\{0, 1\}^i$ with k colors. Our goal is to show that if i is sufficiently large, then with a large probability $g_i \circ \pi_i$ is (r, ϵ) -quasirandom.

Claim 1 *For every $\epsilon_1 > 0$ and $n \in \mathbb{N}$, there is an index j such that if $x_1, \dots, x_n \in X$ are independent ν -random points, then with probability $1 - \epsilon_1$, the map π_j separates all the points in $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$.*

It is easy to see that $\pi = (\pi_1, \pi_2, \dots)$ separates the points of $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$ with probability one on X^n . Let Y_j denote the set of points (x_1, x_2, \dots, x_n) in X^n for which π_j separates the points in $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$. We have that Y_j is an increasing chain of measurable sets such that $\nu(\cup_{i=1}^\infty Y_i) = 1$. This shows that for some index j , we have $\nu(Y_j) > 1 - \epsilon_1$ and finishes the proof of Claim 1.

Let $x = (x_1, \dots, x_n) \in X^n$ and let g be a k -coloring of $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$. Let us say that x is *representative* if the distribution of the l -colored neighborhood $N_{\mathcal{G},h,r}(x_t)$ for a random $t \in [n]$ is $\epsilon/6$ -close to the distribution $\mu_{r,h} := P_{\mathcal{G},r}[h]$. Let us say that (x, g) is *representative* if the distribution of the kl -colored neighborhood $(N_{\mathcal{G},h,r}(x_t), g)$ is $\epsilon/3$ -close to the distribution $\mu_{r,h,k}$.

Let $x = (x_1, \dots, x_n) \in X^n$ be chosen randomly and independently from the distribution ν . We note that with probability 1, the neighborhoods $N_{\mathcal{G},r}(x_i)$ are disjoint. If n is large enough, then (just by the Law of Large Numbers)

$$\Pr_x(x \text{ representative}) \geq 1 - \frac{\epsilon}{6}.$$

Hence if g is a uniform random k -coloring of $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$, and n is large enough, then (by the Law of Large Numbers again), we have

$$\Pr_{x,g}((x, g) \text{ representative}) \geq 1 - \frac{\epsilon}{3}.$$

Let us fix n so that this holds.

Next, using Claim 1, we fix j so that (for a random x) π_j separates all the points in $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$ with probability at least $1 - \epsilon/3$. Whenever this happens, the restriction of $g_j \circ \pi_j$ to $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$ is a uniform random k -coloring. In other words, we can generate a uniform random k -coloring of $\cup_{i=1}^n N_{\mathcal{G},r}(x_i)$ by restricting $g_j \circ \pi_j$ to it if π_j separates it, and randomly k -coloring it otherwise. Thus

$$\Pr_{x,g_j}((x, g_j \circ \pi_j) \text{ representative}) \geq \Pr_{x,g}((x, g) \text{ representative}) - \frac{\epsilon}{3} \geq 1 - \frac{2\epsilon}{3}.$$

It follows that there is at least one k -coloring g_j for which

$$\Pr_x((x, g_j \circ \pi_j) \text{ representative}) \geq 1 - \frac{2\epsilon}{3}.$$

Let us fix such a g_j . Then $c = g_j \circ \pi_j$ is an (r, ϵ) -quasirandom k -coloring of X . In fact, we can generate a random point of x by first generating n independent random points x_1, \dots, x_n and choosing one of them, x_t , uniformly at random. Then with probability at least $1 - 2\epsilon/3$, $(x, g_j \circ \pi_j)$ is representative, and whenever this happens, the distribution of the kl -colored neighborhood $(N_{\mathcal{G},h,r}(x_t), g_j \circ \pi_j)$ is $\epsilon/3$ -close to the distribution $\mu_{r,h,k}$. It follows that the total variation distance of $(N_{\mathcal{G},h,r}(x_t), g_j \circ \pi_j)$ from $\mu_{r,h,k}$, when x_t is also randomly chosen, is at most ϵ . \blacksquare

Our next lemma shows that we can approximate any measurable k -coloring of \mathcal{G}^+ by a k -coloring that is local in the sense of Definition 4.2, and depends only on a discrete approximation of the nodeweights. To be precise, we define the (m, s) -discretization ($m, s \in \mathbb{N}$) as the map $\xi_{m,s} : X^+ \rightarrow U^{s,m}$, where $\xi_{m,s}(x)$ is obtained by considering the neighborhood $N_{\mathcal{G}^+,s}(x)$, and replacing every nodeweight $\xi(v)$ by $\lceil m\xi(v) \rceil$. Recall that the local isomorphism $\phi : V(\mathcal{G}^+) \rightarrow V(\mathcal{G})$ is defined as $\phi : (x, \xi) \mapsto x$.

Lemma 7.10 *For every $r \geq 1$ and $\varepsilon > 0$, and every measurable k -coloring c of \mathcal{G}^+ , there are positive integers s, m and l , a measurable l -coloring h of \mathcal{G} , and a map $f : U^{s,m} \times [l] \rightarrow [k]$ such that the k -coloring $c'(x) = f(\xi_{s,m}(x), h(\phi(x)))$ of \mathcal{G}^+ satisfies*

$$d_{\text{var}}(P_{\mathcal{G}^+,r}[c], P_{\mathcal{G}^+,r}[c']) \leq \varepsilon.$$

Proof. Let (X^+, ν^+) be the underlying space of \mathcal{G}^+ . Let \mathcal{K} denote the set of all subsets of X^+ of the form $\xi_{m,s}^{-1}(y) \cap \phi^{-1}(B)$, where $y \in U^{s,m}$, and B is a Borel set of X . These sets generate the Borel sets of X^+ , hence by the Monotone Class Theorem, the closure under pointwise convergence of the vector space generated by their indicator functions contains every bounded Borel function on X^+ .

In particular, there are pairs of integers (m_i, s_i) , colored balls $y_i \in U^{s_i, m_i}$, Borel sets $B_i \subseteq X$ and real coefficients a_i ($i = 1, \dots, N$) such that

$$\nu^+ \left\{ x \in X^+ : \left| c(x) - \sum_{i=1}^N a_i \mathbb{1}(\xi_{m_i, s_i}(x) = y_i, \phi(x) \in B_i) \right| \geq \frac{1}{2} \right\} < \frac{\varepsilon}{d^{r+1}}.$$

Let $s = \max_i s_i$, $m = \prod_i m_i$, $l = 2^N$, and let h be a Borel l -coloring of X in which every B_i is a union of color classes. Then the sum in the above expression can be written as $g(\xi_{s,m}(x), h(\phi(x)))$ for some $g : U^{s,m} \times [l] \rightarrow \mathbb{R}$. Rounding the values of g to the closest integer in $[k]$, we get a k -coloring c' for which

$$\nu \{ x \in X^+ : c(x) \neq c'(x) \} < \frac{\varepsilon}{d^{r+1}}.$$

For a random point $x \in X^+$, the probability that the colorings c and c' differ on any node in its r -neighborhood is less than ε . This implies the lemma. \blacksquare

Now we are able to prove the main theorem in this section.

Proof of Theorem 7.6. Our goal is to approximate every element in $Q_{\mathcal{G}^+,r,k}$ by an element in $Q_{\mathcal{G},r,k}$ with arbitrary precision $\varepsilon > 0$. In other words, we want to construct, for every measurable k -coloring c of \mathcal{G}^+ , a measurable k -coloring c_0 of \mathcal{G} that defines a similar distribution of colored neighborhoods.

By Lemma 7.10 we may assume that c is of the form $f(\xi_{s,m}(x), h(\phi(x)))$ where h is an l -coloring of \mathcal{G} and $f : U^{s,m} \rightarrow [k]$. Let q be an (s, ε) -quasirandom m -coloring of (\mathcal{G}, h) guaranteed by Lemma 7.9, and let $\mathcal{G}' = (\mathcal{G}, h \times q)$. Consider the k -coloring of \mathcal{G} defined by $c_0(z) = f(N_{\mathcal{G}',s}(z), h(z))$. We claim that c_0 has similar statistics as c :

$$d_{\text{var}}(P_{\mathcal{G}^+,r}[c], P_{\mathcal{G},r}[c_0]) \leq \varepsilon.$$

This follows if we prove that the distributions of $(\xi_{s,m}(y), h(\phi(y)))$ (where y is a random point of \mathcal{G}^+) and $(N_{\mathcal{G}',s}(x), h(x))$ (where x is a random point of \mathcal{G}) are close. But the distribution of

$(\xi_{s,m}(y), h(\phi(y)))$ is just $\mu_{s,h,m}$, and the distribution of $(N_{\mathcal{G}',s}(x), h(x))$ is ε -close to this by the quasirandomness of q . This completes the proof. \blacksquare

The following fact shows another connection between a graphing and its associated Bernoulli graphing. We say that two graphings are *bi-locally isomorphic* if there exists a third graphing that has local isomorphisms into both. The construction of the Bernoulli lift implies that every graphing is bi-locally isomorphic to its Bernoulli graphing. Since by the definition of the Bernoulli graphing, two graphings are locally equivalent if and only if they have the same Bernoulli graphing, we get the following more explicit characterization:

Proposition 7.11 *Two graphings are locally equivalent if and only if they are bi-locally isomorphic.*

To prove this proposition, it suffices to show that bi-local isomorphism is a transitive relation. This takes some work which we do not discuss here; for the details, we refer the reader to [20].

Let us turn to graph sequences. Theorem 7.6 motivates the next definition.

Definition 7.12 (Bernoulli graph sequences) A graph sequence is called Bernoulli if it converges to a Bernoulli graphing in the local-global sense.

Note that the Bernoulli graphing to which a given Bernoulli sequence converges is fully determined by the local limit of the sequence. Bernoulli sequences are basically those which have the least possible global structure among sequences with the same local limit. The following provocative conjecture was popularized by the third author in the past few years. The considerable effort put into the topic shows that the solution may require a substantial novel idea.

Conjecture 7.13 (Limits of random d -regular graphs) Let d be a fixed natural number and let G_i be a random d regular graph on i vertices (if d is odd, then we only consider even values of i). Then $(G_i)_{i=1}^\infty$ is a Bernoulli sequence with probability one.

The motivation for this conjecture is that one expects randomness destroys as much global structure as possible. Notice that if Conjecture 7.13 is true then the limit object is the Bernoulli graphing produced from the d -regular tree. Even the next two weaker conjectures are unsolved.

Conjecture 7.14 A growing sequence of random d -regular graphs is local-global convergent with probability one.

Conjecture 7.15 For every d , there is a Bernoulli graph sequence $(G_i)_{i=1}^\infty$ whose local limit is the d -regular tree.

8 Non-standard graphings

Let $(G_i)_{i=1}^\infty$ be an arbitrary graph sequence of maximum degree at most d . Let ω be a non-principal ultrafilter on \mathbb{N} . Let \mathbf{G} denote the ultraproduct of the graph sequence. The vertex set \mathbf{V} of \mathbf{G} is the ultraproduct of the vertex sets V_i of G_i and the edge set $\mathbf{E} \subset \mathbf{V} \times \mathbf{V}$ is the ultra product of the edge sets $E_i \subset V_i \times V_i$ of G_i . The graph \mathbf{G} has maximum degree at most d , since this property is

expressible by a first order formula. We can also construct a σ -algebra \mathcal{A} on \mathbf{V} and a probability measure μ on \mathbf{V} which is the ultralimit of the uniform distributions on the sets V_i . It is not hard to check that \mathbf{G} satisfies the graphing axiom (1).

If $(G_i)_{i=1}^\infty$ is a locally convergent graph sequence, then \mathbf{G} has neighborhood frequencies that are the limits of the neighborhood frequencies of the graphs G_i . If $(G_i)_{i=1}^\infty$ is locally-globally convergent, then $Q_{\mathbf{G},r,k}$ is the Hausdorff limit of the sets $Q_{G_i,r,k}$.

However, this does not directly prove Theorem 3.2, since (\mathbf{V}, μ) is not a separable probability space. One can complete the proof by choosing an appropriate separable sub-sigma-algebra of \mathbf{G} which preserves the graphing structure. We omit the details here.

An attractive feature of ultralimit graphings is that the sets $Q_{\mathbf{G},r,k}$ are all closed. It is not clear if there is a standard graphing representation of the limit of a convergent sequence with this stronger property.

Question 8.1 Let $(G_n)_{n=1}^\infty$ be a local-global convergent sequence of graphs. Is there a graphing \mathcal{G} that represents the limit with the property that $Q_{\mathcal{G},r,k}$ are all closed?

9 Hyperfinite graphs and graphings

For a graph G , we define $\tau_q(G)$ as the smallest t such that deleting t appropriate nodes, every connected component of the remaining graph has at most q nodes. We say that a sequence $(G_n)_{n=1}^\infty$ of finite graphs is (q, ε) -hyperfinite if $\liminf_n \tau_q(G_n)/|V(G_n)| \leq \varepsilon$. We say that $(G_n)_{n=1}^\infty$ is hyperfinite if for every $\varepsilon > 0$, there is a q such that $(G_n)_{n=1}^\infty$ is (q, ε) -hyperfinite. We can define hyperfiniteness of a graphing \mathcal{G} on underlying space X similarly: let $\tau_q(\mathcal{G})$ denote the infimum of numbers $\delta \geq 0$ such that we can delete a Borel set $S \subseteq X$ with measure δ so that every connected component of the remaining graphing has at most q nodes. We say that a graphing \mathcal{G} is (q, ε) -hyperfinite if $\tau_q(\mathcal{G}) \leq \varepsilon$, and we say that \mathcal{G} is hyperfinite if for every $\varepsilon > 0$, there is a q such that \mathcal{G} is (q, ε) -hyperfinite. Since we are talking about graphs with bounded degree, we could replace deleting nodes by deleting edges in the definitions of hyperfiniteness.

Hyperfiniteness in different settings was introduced by different people (see Kechris and Miller [18], Elek [15], Schramm [23]). Schramm proved that a locally convergent sequence of graphs is hyperfinite if and only if its limit is hyperfinite. This does not hold for (q, ε) -hyperfiniteness for a fixed pair q and ε . As an easy example, a sequence of random d -regular graphs tend to a limiting involution-invariant distribution (concentrated on the infinite d -regular tree) that is $(1, 1/2)$ -hyperfinite, while the sequence is not. On the other hand, a local-global convergent sequence of graphs behaves nicer:

Proposition 9.1 *Let a sequence $(G_n)_{n=1}^\infty$ of finite graphs converge to a graphing \mathcal{G} in the local-global sense. Then $(G_n)_{n=1}^\infty$ is (q, ε) -hyperfinite if and only if \mathcal{G} is (q, ε) -hyperfinite.*

Proof. A finite graph G satisfies $\tau_q(G) \leq \varepsilon |V(G)|$ if and only if it has a 2-coloring c such that $P_{G,k,r}[c](c(\text{root}) = 1) \leq \varepsilon$ and $P_{G,k,r}[c](B) = 0$ for every colored r -ball B that contains a connected all-blue subgraph with $k+1$ nodes. A graphing \mathcal{G} satisfies $\tau_q(\mathcal{G}) \leq \varepsilon$ if and only if for every $\varepsilon' > \varepsilon$, it has a 2-coloring c such that $P_{\mathcal{G},k,r}[c](c(\text{root}) = 1) \leq \varepsilon'$ and $P_{\mathcal{G},k,r}[c](B) = 0$ for every colored r -ball

B that contains a connected all-blue subgraph with $k + 1$ nodes. The proposition follows by the definition of local-global convergence to a graphing. \blacksquare

The following important property of hyperfiniteness is closely related to the results of Schramm [23] and Benjamini, Schramm and Shapira [7]. It can be derived using the graph partitioning algorithm of Hassidim, Kelner, Nguyen and Onak [13]; a direct proof is given in [20].

Proposition 9.2 *Hyperfiniteness is invariant under local equivalence.*

Together with Proposition 9.1, this implies the above mentioned result of Schramm that a locally convergent sequence of graphs is hyperfinite if and only if its limit is hyperfinite. We note that (q, ε) -hyperfiniteness for a fixed q and ε is not invariant under local equivalence, which is shown, for example, by the local-global limits of random d -regular graphs and of random d -regular bipartite graphs. Our main result about hyperfinite graphings is a strengthening of Corollary 7.7.

Theorem 9.3 *Every atom-free hyperfinite graphing \mathcal{G} is local-global equivalent to its Bernoulli graphing.*

Proof. By Corollary 7.7, $\mathcal{B}_{\mathcal{G}} \prec \mathcal{G}$. It remains to show that $\mathcal{G} \prec \mathcal{B}_{\mathcal{G}}$. In other words, for every coloring of \mathcal{G} , we have to find a coloring of $\mathcal{B}_{\mathcal{G}}$ with almost the same local statistics.

Let (X, ν) be the underlying space of \mathcal{G} , let $c : X \rightarrow [k]$ be a measurable coloring, and let us fix a radius $r \in \mathbb{N}$ and an $\varepsilon > 0$. Let ν_B denote the measure of $\mathcal{B}_{\mathcal{G}}$, and set $\varepsilon_1 = \varepsilon / (8(d+1)^r)$. By Proposition 9.2, the Bernoulli graphing $\mathcal{B}_{\mathcal{G}}$ of a hyperfinite graphing \mathcal{G} is also hyperfinite. Let $S \subset \mathfrak{G}[0, 1]$ be a subset such that $\nu_B(S) \leq \varepsilon_1$ and every connected component of $\mathfrak{G}[0, 1] \setminus S$ has at most n nodes. Let $m \in \mathbb{N}$, and define the coloring $b : \mathfrak{G}[0, 1] \rightarrow [m] \times \{0, 1\}$ by $b(x) = (\lceil w(o)m \rceil, \mathbb{1}_S(x))$ where $x = (G, o, w)$. Choosing m large enough, we may assume that the set S' of points x for which $N_{\mathcal{B}_{\mathcal{G}}, r}(x)$ contains two points with the same color has measure at most ε_1 . Note that $\nu_B(S \cup S') \leq 2\varepsilon_1$ and all points of $\mathcal{B}_{\mathcal{G}} \setminus (S \cup S')$ are contained in connected components that have at most n vertices, and whose nodes are colored differently by b .

On the other hand, by Corollary 7.7 we have $\mathcal{B}_{\mathcal{G}} \prec \mathcal{G}$ which implies that there is a coloring $b^* : X \rightarrow [m] \times \{0, 1\}$ such that

$$d_{\text{var}}(P_{\mathcal{G}, n}[b^*], P_{\mathcal{B}_{\mathcal{G}}, n}[b]) \leq \varepsilon_1.$$

It follows that there are subsets $T \subseteq \mathfrak{G}[0, 1]$ and $T' \subseteq X$ with $\nu_B(T) = \nu(T') \leq 4\varepsilon_1$ such that the following conditions hold:

- (a) All points of $\mathcal{B}_{\mathcal{G}} \setminus T$ are contained in connected components that have at most n vertices and whose nodes are colored differently by b , and the same holds for the connected components of $\mathcal{G} \setminus T'$ with coloring b^* ;
- (b) Furthermore, for every $([m] \times \{0, 1\})$ -colored connected graph H with at most n vertices, the measure of points in components isomorphic to H (as colored graphs) is the same in $\mathcal{B}_{\mathcal{G}} \setminus T$ and $\mathcal{G} \setminus T'$. Let $V_H(\mathcal{B}_{\mathcal{G}} \setminus T)$ and $V_H(\mathcal{G} \setminus T')$ be these two sets.

Let C be a connected component of $\mathcal{G} \setminus T'$. Since the vertices of C are colored differently by b^* , there is a (unique) function $f_C : [m] \times \{0, 1\} \rightarrow [k]$ such that $c = f_C \circ b^*$ on the nodes of C . This splits every set $V_H(\mathcal{G} \setminus T')$ into at most k^{2m} measurable sets $V_{H,f}(\mathcal{G} \setminus T')$ (indexed by functions $f : [m] \times \{0, 1\} \rightarrow [k]$) that are unions of components of $\mathcal{G} \setminus T'$.

Split $V_H(\mathcal{B}_G \setminus T)$ into sets $V_{H,f}(\mathcal{B}_G \setminus T)$ so that each $V_{H,f}(\mathcal{B}_G \setminus T)$ is a union of components of $\mathcal{B}_G \setminus T$, and moreover $\nu_B(V_{H,f}(\mathcal{B}_G \setminus T)) = \nu(V_{H,f}(\mathcal{G} \setminus T'))$. This is possible since there is no probability mass on any component of \mathcal{B}_G .

Let c' be the measurable k -coloring of \mathcal{B}_G defined in the following way. Every $v \in V_{H,f}(\mathcal{B}_G \setminus T)$ is colored by $f \circ b(v)$, and the points in T are all colored with one arbitrary color in $[k]$. Note that the (conditional) local statistics of c' obtained by picking a random $v \in \mathcal{B}_G$ conditioned on $N_{\mathcal{B}_G,r}(v) \cap T = \emptyset$ is the same as the (conditional) local statistics of c obtained by picking a random $v \in \mathcal{G}$ conditioned on $N_{\mathcal{G},r}(v) \cap T' = \emptyset$. The ν -measure of the vertices $v \in \mathcal{G}$ with $N_{\mathcal{G},r}(v) \cap T' \neq \emptyset$ is at most $\nu(T')(d+1)^r \leq 4\varepsilon_1(d+1)^r$. The same bound also holds for the ν_B -measure of the vertices $v \in \mathcal{B}_G$ with $N_{\mathcal{B}_G,r}(v) \cap T \neq \emptyset$. Thus we have

$$d_{\text{var}}(P_{\mathcal{G},r}[c], P_{\mathcal{B}_G,r}[c']) \leq 8(d+1)^r \varepsilon_1 \leq \varepsilon$$

which proves the theorem. ■

Remark 9.4 As the proof of Theorem 9.3 shows, $\mathcal{G} \prec \mathcal{B}_G$ holds for every hyperfinite graphing \mathcal{G} (not necessarily atom-free).

Now we are ready to state and prove our main theorem about convergence of hyperfinite graph sequences.

Theorem 9.5 *Every locally convergent hyperfinite graph sequence $(G_n)_{n=1}^\infty$ with $|V(G_n)| \rightarrow \infty$ is a local-global convergent Bernoulli sequence.*

Proof. Let $(G_i)_{i=1}^\infty$ be a locally convergent hyperfinite sequence, and let μ be the involution-invariant measure on \mathfrak{G} that is the local limit of the sequence. Since the Bernoulli graphing \mathcal{B}_μ is locally equivalent to the local limit of $(G_i)_{i=1}^\infty$, Proposition 9.2 implies that it is hyperfinite.

To prove the theorem, assume by contradiction that $(G_i)_{i=1}^\infty$ does not converge in the local-global sense to \mathcal{B}_μ . Then it has a local-global convergent subsequence whose limit graphing \mathcal{G} is not local-global equivalent to \mathcal{B}_μ . By Remark 7.5 the condition $|V(G_n)| \rightarrow \infty$ implies that \mathcal{G} is atom-free. This however contradicts Theorem 9.3. ■

Corollary 9.6 *Local-global convergence is equivalent to local convergence when restricted to growing hyperfinite graph sequences.*

10 Graphings as operators and expander graphings

Let \mathcal{G} be a Borel graph on the probability space (X, μ) with all degrees at most d . If $f : X \rightarrow \mathbb{C}$ is a measurable function, then we define $\mathcal{G}f$ by

$$\mathcal{G}f(x) = \sum_{(x,v) \in E(\mathcal{G})} f(v).$$

It takes a short calculation to show that if \mathcal{G} is a graphing, then it acts on the Hilbert space $L^2(X, \nu)$ as a bounded self-adjoint operator. Let $f : X \rightarrow \mathbb{C}$ be an arbitrary function in $L^2(X, \nu)$. Then

$$\int_x |\mathcal{G}f(x)|^2 d\nu \leq \int_x d \sum_{(x,v) \in E(\mathcal{G})} |f(v)|^2 d\nu = d \int |f(x)|^2 \deg(x) d\nu \leq d^2 \|f\|_2^2.$$

The equality in the above calculation uses the fact that \mathcal{G} satisfies (1). It is easy to see that (1) is equivalent to the statement that the action of \mathcal{G} is self-adjoint in the sense that $\langle \mathcal{G}f, g \rangle = \langle f, \mathcal{G}g \rangle$ holds for every pair f, g of bounded measurable functions. This implies that the action of \mathcal{G} is also self-adjoint on $L^2(X, \nu)$. The Laplace operator corresponding to a graphing is defined as $L = D - \mathcal{G}$ where $Df(x) = f(x)\deg(x)$. It is easy to check that

$$\langle Lf, f \rangle = \int_{(v,w) \in E(\mathcal{G})} (f(v) - f(w))^2 d\eta^* \quad (5)$$

holds in $L^2(X, \nu)$ where η^* is defined in Section 3. Thus L is positive semidefinite .

The theory of graphings is closely related to the theory of measure preserving systems (In a sense, it generalizes ergodic theory.). In particular, one can define the notion of ergodicity. A graphing \mathcal{G} is *ergodic* if there is no measurable partition of the vertex set X into positive measure sets X_1, X_2 such that there is no edge between X_1 and X_2 , or equivalently such that X_1 is a union of connected components of \mathcal{G} . Note that graphings, when defined on an uncountable set, are never connected as graphs and so the notion of ergodicity is a good replacement for the notion of connectivity. Equation (5) implies the following analogue of a well known theorem from ergodic theory about the Koopman representation (see [17]).

Proposition 10.1 *Let L be the Laplace operator corresponding to the graphing \mathcal{G} . The multiplicity of the eigenvalue 0 of L as an operator on $L^2(X, \nu)$ is 1 if and only if \mathcal{G} is ergodic.*

Graphings offer new phenomena. Ergodicity is equivalent to saying that $\nu(N_1(S)) > \nu(S)$ for every set S with $0 < \nu(S) \leq 1/2$ (Here $N_1(S) = \cup_{x \in S} N_{\mathcal{G},1}(x)$). Positive expansion is a natural strengthening of this condition. Let us restrict our attention to d -regular graphs and graphings. We say that a graphing \mathcal{G} is a *c-expander* if for every Borel set $S \subseteq X$ with $0 < \nu(S) \leq 1/2$, we have $\nu(N_1(S)) \geq (1 + c)\nu(S)$. We say that a graphing \mathcal{G} is an *expander* if it is a *c-expander* for some $c > 0$.

Let $(G_n)_{n=1}^\infty$ be a sequence of d -regular graphs that are expanders with expansion $c > 0$. Let us select a local-global convergent subsequence. Then its limit is a d -regular graphing that is also a *c-expander*.

We can generalize spectral conditions for expanders to graphings. Let us define *spectral gap* of a d -regular graphing by

$$\text{gap}(\mathcal{G}) = \inf\{\langle Lf, f \rangle : \langle f, f \rangle = 1, \langle f, 1 \rangle = 0\}.$$

Note that in the above formula the infimum is the same on $L^2(X)$ and $L^\infty(X)$. The following analogue of the theorems of Alon and Milman [4] and Alon [2] on expanders can be proved along the same lines:

Proposition 10.2 *Suppose that a d -regular graphing \mathcal{G} is a c -expander. Then $c^2/(2d) \leq \text{gap}(\mathcal{G}) \leq 2c$. In particular, a graphing is an expander if and only if its spectral gap is positive.*

An easy calculation shows that if \mathcal{G}_1 and \mathcal{G}_2 are local-global equivalent, then $\text{gap}(\mathcal{G}_1) = \text{gap}(\mathcal{G}_2)$. In other words $\text{gap}(\mathcal{G})$ is a local-global invariant quantity. This follows from the classical fact that measurable functions can be arbitrarily well approximated by step functions.

One must be careful though: the spectral gap $\text{gap}(\mathcal{G})$ is a lower bound on the eigenvalues of \mathcal{G} belonging to non-constant eigenfunctions of \mathcal{G} , but it may not be the infimum of such eigenvalues. For example, the Bernoulli graphing of a 2-way infinite path is ergodic but not an expander, and its Laplacian has no non-constant eigenfunction.

11 Graphings and local algorithms

Elek and Lippner [12] formulate a correspondence between graphings and local algorithms. We can make this more precise using the notions of Bernoulli graphings and factor of i.i.d. processes:

Measurable graph theoretic statements for Bernoulli graphings correspond to randomized local algorithms for finite graphs.

Let us start with an example. Let T be the d -regular tree with a distinguished root and let Ω be the compact space $[0, 1]^{V(T)}$. Let $f : \Omega \rightarrow [k]$ be any measurable function which depends only on the isomorphism class of the labeled rooted tree. In other words f is invariant under the action of the root preserving automorphism group of T . Using the function f , we create a random model of k colorings of T in the following way. First we produce a random element $\omega \in \Omega$ by putting independent random elements from $[0, 1]$ on the vertices of T , and then for every $v \in V(T)$, we define the color $c(v)$ as the value of f on the labeled rooted tree obtained from T by assigning labels ω and placing the root on v . We say that f is the *rule* of the coloring process c . Such processes on the tree are called *factor of i.i.d. processes*. We say that the rule f has radius r if it depends only on the labels on vertices of T that are of distance at most r from the root.

The following rule (of radius one) is a classical method to construct an independent set of nodes in a graph (see Alon and Spencer [5]). Let $f : \Omega \rightarrow \{0, 1\}$ be the function which returns 1 if and only if the label on the root is smaller than the labels on all the neighboring vertices. It is clear that with probability one the corresponding random coloring c is the characteristic function of some independent set on T . We can view c as a randomized algorithm which produces an independent set of points of density $1/(d+1)$. Since the rule f has radius one, it can also be applied to a finite d -regular graph G . Let us put random labels from $[0, 1]$ on the vertices of G , and then evaluate the rule f at each vertex using only the neighborhood of radius 1. We get a random $\{0, 1\}$ coloring of $V(G)$ such that 1's form an independent set. Such algorithms (corresponding to a rule of bounded radius) are called local algorithms. On the other hand, we can view f as the characteristic function of a single (non-random) independent set in the Bernoulli graphing \mathcal{G} corresponding to the tree T (That is, $\mathcal{G} := \mathcal{B}_\mu$ where μ is the dirac probability measure on the point $T \in \mathfrak{G}$). The vertex set of \mathcal{G} is $\mathfrak{G}[0, 1]$, but in \mathcal{G} almost every vertex is represented by an element in $[0, 1]^{V(T)}$, and so we

can evaluate the function f for almost every point. It is clear now that $f^{-1}(1)$ is an independent measurable set in \mathcal{G} .

A general definition of factor of i.i.d. processes can be obtained through Bernoulli graphings. Let μ be an involution-invariant measure on \mathfrak{G} , and let \mathcal{B}_μ be the corresponding Bernoulli graphing on $\mathfrak{G}[0, 1]$. Let $f : \mathfrak{G}[0, 1] \rightarrow [k]$ be a Borel function. Then the involution-invariant measure $\mu_{\mathcal{B}, f}$ on $\mathfrak{G}[k]$ has the property that it projects to μ when the labels on the vertices are forgotten. In other words $\mu_{\mathcal{B}, f}$ puts a k -coloring process on the graphs generated by μ . The measure $\mu_{\mathcal{B}, f}$ is called a factor of i.i.d. process on μ . The rule of the process is the function f . We say that the rule f has radius r if $f(G_1) = f(G_2)$ whenever the balls of radius r in G_1 and G_2 are isomorphic as rooted labeled graphs.

We can approximate the rule f with an arbitrary precision ϵ with another rule f' of finite radius r (which depends on ϵ) in the sense that $\nu(x|f(x) \neq f'(x)) \leq \epsilon$. An advantage of the finite radius approximation is that it can be used for local algorithms on finite graphs. Let G be a finite graph of maximal degree at most d , and let us put random labels from $[0, 1]$ on the vertices in G . Then f' defines a new coloring of G such that the color of a vertex v is computed using f' for the labeled neighborhood of radius r of v .

Nondeterministic property testing. The connection between the two convergence notions can be illuminated by the following algorithmic considerations. Given a (very large) graph G with bounded degree, we use the following sampling method to gain information: we select randomly and uniformly a node of G , and explore its neighborhood of radius r . We can repeat this t times. There are a number of algorithmic tasks (parameter estimation, property testing) that can be studied in this framework; we only sketch a simple version of property testing, and its connection to local-global convergence.

It will be convenient to introduce the edit distance for graphs with bounded degree. For two graphs on the same node set $V(G) = V(G')$, we define

$$d_1(G, G') = \frac{1}{|V(G)|} |E(G) \Delta E(G')|.$$

For a graph property \mathcal{P} , let $\mathcal{P}_{-\epsilon} = \{G \in \mathcal{G} : d_1(G, \mathcal{P}) > \epsilon\}$.

We say that the graph property \mathcal{P} is *testable* if for every $\epsilon > 0$, there are integers $r, t \geq 1$ such that given any graph G that is large enough, taking t samples of radius r as described above, we can guess whether the graph has property \mathcal{P} : if $G \in \mathcal{P}$, then our guess should be “YES” with probability at least $2/3$; if $G \in \mathcal{P}_{-\epsilon}$, then the answer should be “NO” with probability at least $2/3$. If \mathcal{P} is *testable*, then a locally convergent graph sequence cannot contain infinitely many graphs from both \mathcal{P} and $\mathcal{P}_{-\epsilon}$.

Now let us say that \mathcal{P} is *nondeterministically testable* if there is an integer $k \geq 1$, and a testable property \mathcal{Q} of k -colored graphs with bounded degree, such that $G \in \mathcal{P}$ if and only if there is a k -coloring c such that $(G, c) \in \mathcal{Q}$. This k -coloring is a “witness” for our conclusion. As an example, the property “ G is the disjoint union of two graphs with at least $|V(G)|/1000$ nodes” is not testable, but it is nondeterministically testable (a witness is a 2-coloring with no edge between the 2 colors). If \mathcal{P} is *nondeterministically testable*, then a local-global convergent graph sequence cannot contain infinitely many graphs from both \mathcal{P} and $\mathcal{P}_{-\epsilon}$.

12 Concluding remarks

Local-global equivalence and limit representation. We have seen a characterization of local equivalence of two graphings (Proposition 7.11). Is there a similar characterization of local-global equivalence?

Does every graphing represent the limit of a local-global convergent graph sequence? This is stronger than the Aldous–Lyons conjecture, but perhaps there is a counterexample. We can mention two possible counterexamples suggested by our results.

Can a d -regular graphing be a better expander than any finite d -regular graph? Such a graphing would certainly be a counterexample. It is not easy, however, to compute the expansion rate of even very simple graphings, like the Bernoulli tree.

Is every graphing $(d+1)$ -edge-colorable in a Borel way? If a graphing is the local-global limit of a sequence of finite graphs, then these graphs can be $(d+1)$ -edge-colored by Vizing’s Theorem, and it is not hard to see that such an edge-coloring can be transferred to the limit graphing.

Even finer limit notions. Limit graphings can represent even finer information than local-global convergence. Consider the following examples. Let $0 < a < 1$ be an irrational number, and consider the following three graphings: (a) \mathcal{C}_a is obtained by connecting every point $x \in [0, 1]$ to the two points $x \pm a \pmod{1}$; (b) \mathcal{C}'_a consists of two disjoint copies of \mathcal{C}_a (both with measure $1/2$); (c) \mathcal{C}''_a is obtained by taking two copies of $[0, 1]$ (call them upper and lower), each with mass $1/2$, and connecting every lower point $x \in [0, 1]$ to the two upper points $x \pm a \pmod{1}$.

These three graphings are locally isomorphic, and either one of them represents the local-global limit of the sequence of cycles. But they are “different”: there is no measure preserving isomorphism between them, and this has combinatorial reasons. The graphing \mathcal{C}'_a is “disconnected” (non-ergodic), while \mathcal{C}''_a is “bipartite”: it has a partition into two sets with positive measure such that every edge connects the two classes. The graphing \mathcal{C}_a does not have any partition with either one of these properties (even if we allow an exceptional subset of measure 0). This follows from basic ergodic theory.

It seems that the graphing \mathcal{C}_a should represent the limit of odd cycles, \mathcal{C}'_a should represent the limit of graphs consisting of a pair of odd cycles, while \mathcal{C}''_a should represent the limit of even cycles. A theory of convergence that would explain this example has not been worked out, however.

We know [9] that local convergence is equivalent to right-convergence where the target graph is in a small neighborhood of the looped complete graph with all edge-weights 1. Can local-global convergence be characterized by some stronger form of right convergence?

References

- [1] D. Aldous, R. Lyons: Processes on Unimodular Random Networks, *Electron. J. Probab.* **12**, Paper 54 (2007), 1454–1508.
- [2] N. Alon: Eigenvalues and expanders, *Combinatorica* **6** (1986), 83–96.

- [3] M. Abert, B. Weiss: Bernoulli actions are weakly contained in any free action, *Ergodic theory and dynamical systems*, to appear.
- [4] N. Alon and V. D. Milman: λ_1 , isoperimetric inequalities for graphs and superconcentrators, *J. Combinatorial Theory B* **38** (1985), 73–88.
- [5] N. Alon, J. Spencer: *The Probabilistic Method*, Wiley–Interscience, 2000.
- [6] I. Benjamini and O. Schramm: Recurrence of distributional limits of finite planar graphs, *Electronic J. Probab.* **6** (2001), paper no. 23, 1–13.
- [7] I. Benjamini, O. Schramm, A. Shapira: Every Minor-Closed Property of Sparse Graphs is Testable, *Advances in Math.* **223** (2010), 2200–2218.
- [8] B. Bollobas, O. Riordan: Sparse graphs: Metrics and random models, *Rand. Struc. Alg.* **39**, 1–38.
- [9] C. Borgs, J.T. Chayes, J. Kahn, L. Lovász: Left and right convergence of graphs with bounded degree, *Random Struc. Alg.*
<http://arxiv.org/abs/1002.0115>
- [10] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi: Convergent graph sequences I: Subgraph frequencies, metric properties, and testing, *Advances in Math.* **219** (2008), 1801–1851.
- [11] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi: Convergent graph sequences II: Multiway Cuts and Statistical Physics, *Annals of Math.* (to appear),
<http://www.cs.elte.hu/~lovasz/ConvRight.pdf>
- [12] G. Elek and G. Lippner: Borel oracles. An analytical approach to constant-time algorithms, *Proc. Amer. Math. Soc.* **138** (2010), 2939–2947.
- [13] A. Hassidim, J.A. Kelner, H.N. Nguyen, K. Onak: Local Graph Partitions for Approximation and Testing, in: *FOCS '09, 50th Ann. IEEE Symp. on Found. Comp. Science* (2009), 22–31.
- [14] G. Elek: On limits of finite graphs, *Combinatorica* **27** (2007), 503–507.
- [15] G. Elek: The combinatorial cost, *Enseign. Math.* bf53 (2007), 225–235.
- [16] G. Elek, B. Szegedy: A measure-theoretic approach to the theory of dense hypergraphs, arXiv:0810.4062
- [17] Eli Glasner: *Ergodic Theory via Joinings*, Amer. Math. Soc. (2003).)
- [18] A. Kechris and B.D. Miller: *Topics in orbit equivalence theory*, Lecture Notes in Mathematics **1852**. Springer-Verlag, Berlin, 2004.
- [19] A. Kechris, S. Solecki and S. Todorcevic: Borel chromatic numbers, *Advances in Mathematics* **141** (1999) 1–44.

- [20] L. Lovász: *Large graphs, graph homomorphisms and graph limits* (forthcoming book).
- [21] L. Lovász, B. Szegedy: Limits of dense graph sequences, *J. Comb. Theory B* **96** (2006), 933–957.
- [22] L. Lovász, B. Szegedy: Szemerédi’s Lemma for the analyst, *Geom. Func. Anal.* **17** (2007), 252–270.
- [23] O. Schramm: Hyperfinite graph limits, *Elect. Res. Announce. Math. Sci.* **15** (2008), 17–23.